

Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, DC

WC Docket No. 17-108
GN Docket No. 14-28

In the Matter of

Restoring Internet Freedom

Protecting and Promoting the Open Internet

Light-Touch Regulation by
Banning Unreasonable Discrimination

Jon M. Peha

Professor, Carnegie Mellon University

Address: Carnegie Mellon University
Department of EPP
Pittsburgh, PA 15213-3890
peha@cmu.edu
www.ece.cmu.edu/~peha

July 17, 2017

Executive Summary

Since the early 2000s, all seven chairmen of the Federal Communications Commission have publicly supported the “Internet freedoms,” which give users the right to access the lawful content, application, and devices of their choice. The public interest is best served by a regulatory environment that gives Internet access providers as much freedom as possible to innovate and expand, while preventing Internet access providers from becoming information gatekeepers or violating these freedoms. Providers of broadband Internet access services have the technical capability and the financial incentive to violate these freedoms. Thus, there should be some *ex ante* Open Internet rules, and only the Federal Communications Commission has the mandate and expertise to establish and enforce such rules. However, it is possible to improve upon the 2015 Open Internet rules. Some portions of the existing rules create too much regulatory uncertainty, such as the “Internet conduct standard.” Other portions do not focus as specifically as they could on Internet service provider strategies that could lead to gatekeeping. We propose an alternate approach to “light-touch” regulation based on the idea that the FCC should prohibit “unreasonable discrimination,” while allowing reasonable discrimination. Discrimination is unreasonable when it is based on content, application, service, or attached device. To be effective, rules must appropriately address specialized services as well. The paper does explain how the proposed rules might fit under either Title I or Title II, but the infamous question of whether the FCC should reclassify is outside our scope. (For discussion of that issue, see [PE17] and [CP14].)

Author Qualifications

Jon Peha is a Full Professor at Carnegie Mellon University, with experience in industry, government, and academia. In government, he served at the FCC as Chief Technologist, in the White House as Assistant Director of OSTP, in the House Energy & Commerce Committee, and at USAID for the Telecommunications Leadership Program. In industry, he has been Chief Technical Officer for three high-tech companies, and member of technical staff at SRI International, AT&T Bell Labs, and Microsoft. At Carnegie Mellon, he is a Professor in the Dept. of Electrical & Computer Engineering and the Dept. of Engineering & Public Policy, and former Associate Director of the Center for Wireless & Broadband Networking. Dr. Peha holds a PhD in electrical engineering from Stanford. He is an *IEEE Fellow* and an *AAAS Fellow*, and was selected by AAAS as one of 40 Featured Science and Technology Policy Fellows of the last 40 years (“40@40”). Dr. Peha has received the FCC’s “Excellence in Engineering Award,” the IEEE Communications Society TCCN Publication Award for career contributions to spectrum management, and the Brown Engineering Medal. He consults on a wide range of technical and policy issues related to information and communications technology.

Disclaimer

In writing this comment, Dr. Peha represents no one but himself.

Apology

To those who have read my previous writings, I apologize for recycling text. But as my mother used to tell me as a child, “if you keep asking me the same question, you’re going to get the same answer.”

1 Why we need some ex ante Open Internet rules enforced by the FCC

I applaud the Commission for proclaiming that (i) we need rules to protect the free and open Internet, and (ii) those rules should not be so onerous that they restrict healthy innovation. As stated in Paragraph 71 of the NPRM [FC17], there has been a “long-standing consensus under both Democratic and Republican-led Commissions, represented by the four Internet Freedoms, that consumers should have access to the content, applications, and devices of their choosing as well as meaningful information about their service, all without deterring the investment and innovation that has allowed the Internet to flourish.”

The Commission at the outset “seek(s) comment on whether *ex ante* regulatory intervention in the market is necessary in the broadband context” ([FC17], paragraph 77). Yes, there is need for some ex ante regulations, because in the absence of FCC regulations, Internet access providers today would have both the technical ability and the financial incentive to violate the four freedoms, and to extract oligopoly rents.

We begin with technical capability. Technology such as deep packet inspection and flow classification emerged in the last decade that allow ISPs to effectively determine the content, application, and device associated with each packet stream [PE06]. That ability has been enhanced by a trend in interconnection whereby some large content and application service providers peer directly with BIAS providers, so that the latter can infer information about application, service, and possibly content from where a packet enters the ISP’s network. Once these packets and packet streams have been identified, an ISP can discriminate based on content, application, service or device in a variety of ways. This includes blocking some packet streams, causing some streams to have superior or inferior quality of service, or causing some streams to yield higher or lower costs to users.

A BIAS provider would have financial incentives to use these capabilities in ways that are contrary to the public interest. One obvious example is that a BIAS provider may also provide services that compete with content providers, application service providers, or vendors of devices that attach to the Internet, all of which depend on the ability to exchange traffic over the Internet. The BIAS provider can use its control over the network to put competitors at a disadvantage. This can be done by blocking traffic based on content, application, service or device, but blocking traffic outright is not necessary. For example, by making sure that users observe a lower quality of service when they access a competitor’s service, the BIAS provider’s service can obtain a large competitive advantage. In this case, quality of service might be measured in throughput, latency, packet loss, or other relevant metrics. These practices can increase profits for BIAS providers, but they do so by undermining competition.

A BIAS provider can also extract monopoly or oligopoly rents from content, application, service or device markets even if the BIAS provider does not compete directly in any of these markets [PE06]. For example, a BIAS provider can extract rents from the online market for digitized songs by adding a 10 cent charge to every song carried by the network from any of the online music vendors, even though there is no such charge to carry other files of identical size. This could push the total cost to customers of a song to the highest value that the market will bear, i.e. the monopoly price, even though there are many competing music vendors. Similarly, a BIAS provider could charge for every voice over IP (VOIP) packet carried, even if the cost of carry VOIP packets is negligible, and if the BIAS provider offers no

voice service of its own. The technology to do so already exists. Alternatively, the BIAS provider could ensure that quality of service for VOIP streams or video-on-demand is poor unless a fee is paid, and this fee could be established to extract monopoly or oligopoly rents that are unrelated to the cost of actually providing the service.

A BIAS provider can also benefit by blocking or throttling attempts to access content that is contrary to the interests of the BIAS provider. Some of this content could make the BIAS provider less attractive to consumers, and thereby reduce the BIAS provider's profits. For example, a BIAS provider may degrade quality of service when a user is accessing the website of a competing BIAS provider, or negative reviews written by the BIAS provider's own customers, or articles from technology journalists that rate the BIAS provider's services as inferior to others on the market. The BIAS provider could similarly try to keep negative assessments of the BIAS provider away from its employees. For example, a Canadian ISP was accused of blocking access to a labor union's website during a labor dispute [PE06]. A BIAS provider can even limit political speech. Many ISPs already give large campaign contributions. Why not make sure that access to the preferred political candidate's website is much faster than access to that candidate's opponent?

The Commission raises the question of whether Department of Justice antitrust oversight would suffice (e.g. [FC17], paragraphs 78 and 84), but the preceding examples demonstrate why FCC rules are still needed. The application of antitrust law is more clear when vertically integrated companies use their control over BIAS access to benefit their own content services by undermining competition, but it is far less obvious how to use antitrust laws against BIAS providers that are extracting monopoly or oligopoly rents from markets with many competitors, or BIAS providers that are limiting the speech of labor unions and politicians. Similarly, the "unfair and deceptive practice" standard enforced by the Federal Trade Commission works well if a BIAS provider violates its own promises, but can otherwise be inadequate. There is also the question of expertise. The Federal Communications Commission has staff who understand Internet technology, and this is essential when formulating and enforcing Open Internet policy. At least with current resources and mandates, the Federal Communications Commission is the appropriate organization to address this issue.

Some have argued that even though BIAS providers have incentive to violate the four freedoms, we should not enact ex ante regulations to prevent this if the record of violations in the U.S. is limited. For example, the NPRM asks the following. "Prior to 2015, many large Internet service providers voluntarily abided by the 2010 no-blocking rule in the absence of a regulatory obligation to do so. Do we have reason to think providers would behave differently today if the Commission were to eliminate the no-blocking rule?" ([FC17], paragraph 80) The answer is yes. At a time when the Commission had established network neutrality principles but had not yet initiated a rulemaking, we saw one of the largest BIAS providers (Comcast) blatantly violating those principles by blocking traffic associated with one application, thereby demonstrating the need for rules. It was this action that put the Commission under both Republican and Democratic leadership on the path to Open Internet regulations, ultimately yielding first the 2010 Order. Although there have been intervals since then when no Open internet rules applied, such as immediately after the court struck down the 2010 rules, there wasn't a single day when the Commission let ISPs believe that rules were not forthcoming, or that BIAS providers would ultimately be free to violate Open Internet principles without risk of future repercussions. If the

Commission were ever to remove all Open Internet rules and announce that no new rules would be forthcoming, then we should expect BIAS providers act in a way that maximizes their profit, which would include violating the four freedoms. The argument that ISP behavior in 2014 is representative of what they would do without Open Internet rules is akin to an argument a convict might make in a parole hearing: “I haven’t robbed any banks while I was in prison, so clearly I won’t rob any banks if you let me out.”

The Commission also asks whether “by disclosing to consumers that it is offering a ‘curated internet experience,’ ... an ISP escapes from the ambit of the rules entirely” and the implications of this possibility ([FC17], paragraph 79). It is possible that a service will emerge that provides only curated content, and it is reasonable to ask what rules if any should apply to such a service, but this would not eliminate the need for Open Internet rules which apply to a BIAS that lacks curation. Many consumers want a service that gives them access to all parts of the Internet, without curation (except what is needed for cybersecurity protection). Many owners of Internet infrastructure want to provide an Internet access service without curation, so the ISP can avoid liability for transfers of illegal content. Thus, uncensored Internet access is not going away any time soon.

2 Why we need greater regulatory certainty

There is an inevitable trade-off when establishing rules for the Internet, telephone networks, cellular networks, or any technology that changes from year to year. If the rules are tied too specifically to current technology and services, then those rules can quickly become meaningless or even counterproductive. Some degree of flexibility is needed. On the other hand, if the rules are too vague, then they create uncertainty that can harm service providers and consumers. When rules are too vague, companies may be reluctant to offer beneficial services that regulators would accept, or may invest resources for purposes that are later prohibited. The 2015 rules could be improved by increasing specificity.

2.1 Uncertainty over the “Internet Conduct Standard”

One example of rules that deserve improvement is the rules that would not allow a BIAS provider to “unreasonably interfere with” or “unreasonably disadvantage,” and the so-called “Internet conduct standard.” It may be appropriate for the Commission to seek rules that instruct BIAS providers not to “unreasonably interfere with or unreasonably disadvantage end users’ ability to select, access, and use broadband Internet access service or the lawful Internet content, applications, services, or devices of their choice,” [FC15], but the Commission should provide greater clarity as to what that entails. It is not entirely clear what it means to “interfere with” or “disadvantage,” since this implies that a packet or packet stream has some natural state, and that the ISP then imposes on that natural state; this does not map easily to how networks actually work. Allocating resources to improve the quality of service (QoS) of some traffic streams often has the effect of decreasing QoS of other streams, and this is not inherently bad. That sounds like it could be disadvantaging, but even if it is, under what circumstances is it *unreasonable* disadvantaging? In the absence of guidance, BIAS providers may be deterred from offering services that would not violate regulations and would benefit consumers, or they may not be deterred from offering services that are harmful to consumers. The 2015 rules do offer factors to consider, which is helpful, but do not make sufficiently clear how and when to apply these factors.

The Commission seeks input regarding whether to eliminate the rule ([FC17], e.g. paragraphs 73-75). It is impossible to answer the question of whether the public interest is better served by eliminating this rule or by making this rule more specific without knowing what other changes are to be made to the other Open Internet rules. For example, if the no-blocking rule were eliminated, then it could be very important to explicitly make blocking based on content, application, or device an example of “unreasonable disadvantaging.” However, this would not be necessary if an effective no-blocking rule is kept. I will address these broader questions in Section 4. However, it is clear that leaving the “Internet conduct standard” in its current form is not the best option.

2.2 Uncertainty over Non-BIAS Data Services / Specialized Services / Managed Services

Another example of something that was specified too vaguely in the 2015 Open Internet rules [FC15] is the “non-BIAS data service,” which has previously been called a “specialized service” or a “managed service.” The Commission asks whether it should “adhere to the characteristics of non-broadband Internet access service data services described in the *Title II Order*?” Or “revert to the general concept of non-broadband Internet access service data services discussed in the *Open Internet Order* (and then known as ‘specialized services’)?” ([FC17], paragraph 94). The answer to all of the above is no. The existing case-by-case approach without guidance on how cases will be resolved will produce regulatory uncertainty. It is time for a better approach.

The FCC has not made sufficiently clear what a specialized service (or non BIAS data service or managed service) is, nor have they have made clear what rules apply to such services. According to the 2015 rules, the FCC will “closely monitor their development and use” of these services [FC15] to determine whether “these types of service offerings are undermining investment, innovation, competition, and end-user benefits.”

These rules could create what I have previously called the “specialized services loophole” [PE14, PE15, PE16], violating the “four freedoms,” and even allowing overtly anticompetitive acts. Perhaps through these allegedly non-BIAS services, the broadband provider will provide superior quality of service when transferring video, audio, and e-books to affiliated content providers at a price of 1 cent per GB, and to unaffiliated content providers at a price of 1 million dollars per GB, thereby gaining monopoly or oligopoly control over any upstream service that requires superior quality of service. In the absence of rules against such practices within non-BIAS data services / specialized services / managed services, broadband providers could easily have the ability and incentive to adopt such practices.

The root cause of this loophole is that the FCC has created a legal distinction where no technical distinction exists, as the examples above demonstrate. There is no practical difference between the case where a provider offers both a BIAS and a separate non-BIAS data service over the same shared capacity, and the case where a provider uses 100% of the capacity to offer a BIAS and some specialized non-BIAS services happen to run on top. Thus, regulations that limit how capacity is allocated within a BIAS, but not how capacity is allocated between BIAS and non-BIAS data services, or how capacity is allocated among non-BIAS data services, will inevitably be easy to ignore.

The appropriate policy for specialized services depends on the rules for BIAS services. Thus, I will address the issue of how best to reduce the current uncertainty in Section 6, after consideration of other rules.

3 Why current rules could be revised or replaced, but not eliminated

As described in Section 2, in the absence of Open Internet rules, BIAS providers would have the technical ability and the financial incentive to limit users' access to content, applications, and devices. They do this by using approaches that discriminate based on content, application, or device to affect one of three things: (i) which traffic passes through their network (i.e. blocking), (ii) the quality of service experienced by that traffic, and (iii) the price charged for carrying that traffic. Open Internet address rules should address all three: availability, quality of service, price.

3.1 Blocking Rule

The Commission seeks "comment on whether a codified no-blocking rule is needed to protect such freedoms," i.e. "the freedom to send and receive lawful content and to use and provide applications and services without fear of blocking is essential to the Internet's openness." ([FC17], paragraph 80). For reasons discussed in Section 2, a codified rule is needed. The current no-blocking rule is an appropriate means to achieve this outcome, as long as the "reasonable network management" definition maintains the current exclusion when blocking traffic that is a cybersecurity threat.

3.2 Quality of Service Discrimination / Throttling Rule

As described above, a BIAS providers can limit access to content, applications, and devices in a way that makes it a gatekeeper by making sure that the associated packet streams experience poor quality of service. Open Internet rules should address this, and in the 2015 Order, the "no-throttling" rule is the one bright-line rule that does. Thus, the rule is necessary. Nevertheless, there are other ways to write rules that meet this objective.

A rule that prevented a BIAS provider from offering superior quality of service for those applications that need it would be overly restrictive, and not in the public interest. However, as I understand it, the current no-throttling rule does not have this problem, i.e. the rule does not "prevent providers from offering broadband Internet access service with differentiated prioritization that benefits consumers, nor does it "harm latency-sensitive applications and content" ([FC17], paragraph 83). For example, a BIAS provider might allow a customer to choose between a low-latency class of service, and a standard-latency class of service. The BIAS provider is then employing algorithms that discriminate based on the quality levels chosen by customers, but not based on the application or content itself.

The word "throttle" is not ideal for the same reasons that the word "disadvantage" is problematic, as described in Section 2; it implies that the BIAS provider reduced quality of service from a natural higher level to a lower level, but it is not always clear what that natural level might be. Indeed, given that some applications care more about latency while others care more about throughput, it is not always obvious what "better" is. It could be clearer to say that a BIAS provider should not adjust quality of service

based on content, application, service, or non-harmful device. This can be done by refining the existing rule slightly, or through the approach discussed in Section 4.

3.3 Price Discrimination / No Paid Prioritization Rule

A BIAS provider can also limit access to content, applications, services and devices in a way that makes it a gatekeeper by making sure that the associated packet streams are very expensive to end users. Open Internet rules should address this, and in the 2015 Order, the “paid prioritization” rule is the one bright-line rule that does. Thus, the rule is necessary, although there may be ways to prevent BIAS providers from becoming gatekeepers in ways that are less restrictive to the BIAS providers than the current rule.

Under the right circumstances, consumers benefit if BIAS providers are allowed to prioritize, or more broadly to offer multiple levels of quality of service, because there can be highly valuable applications that require superior quality of service, such as remote health care monitoring. Since there is typically a cost to providing superior quality of service, BIAS providers are far more likely to provide these capabilities if they can charge for superior quality of service. The current paid prioritization rules do allow BIAS providers to charge for superior quality of service, but they may only charge their direct subscribers, and not edge providers. Thus, the rule does not prevent the “development of real-time or interactive services” ([FC17], paragraph 86).

Nevertheless, the risk to an open Internet does not come simply when BIAS providers offer prioritization, or when they charge for prioritization, but when BIAS providers apply charges unfairly in a way that limits access to content, application, service, or device. For example, a BIAS provider may charge \$1 per byte of traffic from a competing content provider. Thus, the no paid prioritization rule could be modified as long as Open Internet rules somehow continue to prohibit unfair price discrimination. One way to do so is discussed in Section 4.

3.4 Reasonable Network Management

The FCC asks whether to maintain the exception for reasonable network management, and more specifically, whether to “eliminate the restriction imposed by the *Title II Order* that the exception will only be considered if used for a ‘technical management justification rather than other business justifications,’ or if we should return to the 2010 definition of ‘reasonable network management’ that did not contain that qualifier.” ([FC17], paragraph 93)

It is not possible to determine the right reasonable network management exception to the Open Internet rules without knowing the rules. In general, the more stringent the Open Internet rules are, the more exemptions are needed. For example, if Open Internet rules prohibit all forms of throttling, then reasonable network management exemptions would be needed that allow a BIAS provider to use some forms of throttling when confronted with certain types of congestion. However, if Open Internet rules only prohibit throttling that discriminates based on content, then the rules might be improved by eliminating any reasonable network management exemption that is related to congestion, because throttling for congestion control should not discriminate based on content.

Nevertheless, I am not currently aware of any changes to the core Open Internet rules that the FCC is considering that should have an exemption for practices that have no technical justification whatsoever, and can only be justified because they are good for the business of the BIAS provider. Such a change in the definition of reasonable network management could probably only be justified as advancing the public interest if the FCC makes the basic Open Internet rules *more stringent* than the 2015 rules.

4 An Alternative: Preventing *Unreasonable* Discrimination

Fundamentally, “network neutrality” and the “Open Internet” and “Internet Freedom” are about discrimination. We should prohibit ISPs from using discrimination to become gatekeepers, but not prohibit more forms of discrimination than is necessary, thereby limiting innovation and deterring investment. The no-blocking, no-throttling and paid-prioritization rules were intended to limit discrimination related to availability, quality of service, and pricing. There should be limits on these three types of discrimination, as discussed in previous sections, but there are advantages to discrimination as well. Consumers can greatly benefit from some uses of discrimination and can be harmed by other uses of discrimination [PE06]. Ideally, we should seek Open Internet rules that encourage the beneficial and prohibit the harmful, although tradeoffs are inevitable.

I propose a simpler approach: prohibiting *unreasonable discrimination*, where discrimination is considered unreasonable when it is based on content, application, service, or attached device. This simple technology-neutral rule protects the first three “Internet Freedoms” first articulated by FCC Chairman Michael Powell, and prevents the most dangerous forms of gatekeeping. In contrast, discrimination based only on the class of service that was selected by the subscriber, or on traffic volume, or on distance to a transmitter, is not inherently unreasonable.

Some applications benefit from superior QoS, which is best provided through some form of discrimination. For example, someday a telemedicine application may allow patients to recover from surgery in their own homes with constant monitoring from medical professionals at a hospital. An ISP could use discrimination to ensure that QoS is adequate for medical monitoring. Thus, some applications that benefit from superior QoS actually involve life and death, although many do not. For example, those recovering patients and their healthy neighbors may want to listen to music streamed by Internet radio stations, but with guarantees that transient congestion will not ruin the sound. To avoid denying Internet users the benefits of these applications, FCC rules should not prohibit the kind of discrimination that makes such applications possible, nor should FCC rules prevent an ISP from charging for these services because the services may not emerge unless someone pays for them. These examples should be considered *reasonable discrimination*, as long as it is the user that chooses the appropriate QoS, and not the ISP that chooses, perhaps over the user’s objections. In the best cases, Internet access providers would simply be allocating limited resources to improve the QoS of those applications that benefit most from those resources, and charging for the resources allocated.

Arrangements like these become problematic when an ISP seeks to extract oligopoly rents through fine-tuned discrimination [PE06]. Perhaps the telemedicine application requires the exact same QoS from the network as a new multiplayer game, but if the network is allowed to set prices based on the application, rather than based on the QoS that the network offers, the ISP will charge very differently for

telemedicine and gaming. With limited competition, the ISP will charge the maximum that the market will bear for each application, and consumers could see higher prices and less availability for content and applications as a result. Alternatively, someone may wish to use the same underlying service that supports the streaming of music to support the streaming of political speeches against mergers in the telecom industry. Would an ISP agree to carry the music, but not carry the speeches? These are examples of discrimination based on application and content, and discrimination by those criteria should be considered *unreasonable discrimination*. If an ISP with market power has unrestricted ability to discriminate based on content and application, then it has the ability and the incentive to give content providers affiliated with the ISP an advantage over competitors. Moreover, it would have incentive to bring separate oligopoly rents onto each distinct application or content market, even when that application or content market is highly competitive, as discussed in Section 2.

In short, it should be considered *unreasonable* to block or discriminate with respect to QoS or price solely based on content, application, user, service, or type of attached device when none of these factors affect the scarce resources that the network must allocate. It should be considered *reasonable* to discriminate between Class A service and Class B service with respect to QoS and price as long as both services are available to all and each user gets to decide between A and B, but *unreasonable* if some users can access one service and other similarly-situated users cannot, or if the ISP decides between A and B rather than the user.

5 Implications for Title I vs. Title II: A Most Favored Nation Approach

The approach described in Section 4 is certainly possible under Title I, if the FCC forbears from imposing many regulations. Although there is uncertainty, it is possible that the approach is possible under Title I, based on Section 706. The DC Appeals Court has indicated that the FCC does have authority to impose some kinds of Open Internet rules under Section 706, and indicated that a policy allowing negotiation is more likely to be allowed. It is possible to prohibit unreasonable discrimination while allowing some degree of negotiation through a “most favored nation” approach I have previously proposed [PE14, PE16].

The FCC could allow negotiation and still prohibit some forms of unreasonable discrimination by allowing content and application service providers to invoke “most favored nation” rights, similar to what we see in international trade. This means that any content or application service provider can demand that an Internet access provider offer it the same deal that this Internet access provider has recently given some other similarly situated content or application service provider. This deal includes both the quality of service offered by the Internet provider, and the pricing. To grant this deal to one content or application service provider and then deny the deal to another similarly situated content or application service provider would be considered unreasonable. Volume discounts should be viewed with skepticism, so that in general, small content and application service providers can get the same deals as large ones. A “most favored nation” rule limits the extent to which an Internet access provider can exploit differences in willingness to pay to extract rents. An Internet access provider could at first serve the upstream markets with high demand at a high price, but once the Internet access provider lowers price to serve other upstream markets with similar quality of service needs, the network would be prevented from raising prices again any time soon.

The Commission could adopt this approach whether it classifies Internet access as a Title I or Title II service. If the Commission chooses to reclassify as Title I, it is unclear whether effective Open Internet rules that protect the four freedoms could be upheld. However, the courts are certainly more likely to find that Section 706 gives the FCC authority for the “most favored nation” approach than for the policy adopted in 2015. The DC Appeals Court has indicated that a policy allowing negotiation similar to cellular roaming is distinct from common carrier regulation [US14], and the “most favored nation” approach allows individual negotiation. Any time an Internet access provider and a content provider wish to adopt an agreement for paid priority that they deem to be of mutual benefit, the parties are entitled to do so, even if the terms of this agreement are unlike any that have come before. These arrangements are therefore made in a manner that is demonstrably different from what is allowed under common carrier regulation, as the Court requires. The “most favored nation” approach also facilitates the virtuous circle of innovation that the court found convincing. Consider that entrepreneur with a new idea for an application. The entrepreneur is free to negotiate directly with the Internet access provider to get better quality of service than is typically available, but this can be costly, and it isn’t necessary. Instead, the entrepreneur could examine the deals that have already been established by knowledgeable negotiators representing well-funded content and application service providers, and try to find the deal that will work best for this new application. This facilitates entry by small innovators.

6 Non-BIAS Data Services / Specialized Services / Managed Services

There are companies that offer both broadband Internet access services and other communications services that are exempt from FCC BIAS regulations. Over the years, these exempt services have been called non-BIAS data services, specialized services, or managed services. In this section, I will refer to them as specialized services, because the phrase “specialized service” brings fewer presumptions (and possibly misconceptions) about what the term might mean.

There is good reason why some communications services should not be subject to BIAS rules, but the existence of these exempt specialized services inherently brings two risks. First, if the definition is not sufficiently clear, then both BIAS providers and regulators may be unsure whether a given service is subject to BIAS rules, or specialized service rules, or both. As discussed in Section 2.2, this uncertainty is a legitimate concern with the 2015 rules. Second, if services that could be viewed as part of a BIAS are instead made exempt from BIAS rules, then there is a risk that BIAS rules may fail to meet their objectives, i.e. we would live in the worst of all worlds where the FCC must create and enforce regulations, but under the guise of “specialized services” BIAS providers are still able to become information gatekeepers, extract oligopoly rents, undermine competition, and interfere with the “four freedoms” of an Open Internet that every Republican and Democratic chairman since Chairman Powell has vowed to support.

In the absence of clear limits on the definition of specialized services, current technology would make it easy for Internet access providers to engage in unreasonable discrimination among their offerings, while maintaining the illusion that their specialized services are somehow distinct from their BIAS. For example, as previously suggested, traffic from a supposed specialized service “may be sent over a separate virtual local area network (VLAN), or a separate service flow in a cable system operating under

the DOCSIS standard” [PE06]. While the VLAN or service flow for a specialized service may appear to be separate from general Internet traffic, traffic from both services travel through the same bottleneck links, and the VLAN or service flow identifier is used to give the supposed specialized service priority over other traffic during periods of congestion. In reality, this is simply a priority Internet service. If priority Internet services receive blanket exemption from Open Internet rules regarding discrimination and blocking because they were given a specialized service label, then BIAS providers can violate the four freedoms at will.

There are two reasonable policy approaches for specialized services. One approach is to establish a relatively narrow definition of specialized services, so that only services that are sufficiently different from Internet access, and in no way compete with Internet access, are exempt from BIAS rules. In this case, there will be little need for rules on the provision of specialized services. The other approach is to create rules for specialized services that protect the four freedoms and prevent the providers of specialized services from becoming gatekeepers. In this case, the definition of specialized service can be much broader without risk. As this section will show, the latter approach is more likely to be best if the FCC imposes more stringent Open Internet rules on Internet access services, and the former is more likely to be best if the FCC imposes less stringent Open Internet rules. Sections 6.1 and 6.2 will describe the two approaches. Both are summarized in Section 6.3.

6.1 Option 1: Heavier Rules on BIAS, Broad Definition of Specialized Service, and Meaningful Regulation of Specialized Services

Some of the services that have been proposed as examples of specialized services are not distinctly different from typical Internet service except that the provider promises superior quality of service. For example, perhaps a telemedicine application requires high quality of service. If Open Internet rules prohibit a BIAS provider from discriminating in order to protect the quality of service for a telemedicine stream, or prevent the BIAS provider from charging for the service, then the BIAS provider would reasonably try to offer this potentially valuable service as a specialized service, and allowing the provider to do so could serve the public interest. In this case, however, there should be light-touch rules in specialized services for the same reasons they are needed in BIAS services, as described in Section 1. For example, the FCC could apply to specialized services the light-touch rules discussed in Section 4, which are based on the concept of “unreasonable discrimination.” As described in Section 2.2, the FCC’s 2015 Order [FC15] indicates that rules could be applied to specialized services in the future, but this policy is not clearly articulated.

6.2 Option 2: Lighter Rules on BIAS, Narrow Definition of Specialized Service, and Little Regulation of Specialized Services

Now consider the case where Open Internet rules allow BIAS providers to offer superior quality of service to users who choose that quality within a BIAS service, and allow BIAS providers to charge accordingly. (I proposed such a policy in Section 4.) In this case, there is little reason to create a “specialized service” category that bypasses BIAS rules for Internet streams with superior quality of service, because BIAS rules are not a significant impediment. Thus, we can exempt a much narrower classes of services under the specialized services label, and if we do this, there is far less need to

regulate specialized services at all. It may even be sufficient to only require basic transparency and truth in advertising, so consumers know what they are buying.

If we consider the case where few rules, if any, are imposed on specialized services, then we must ask whether there are cases where FCC rules should allow ISPs to engage in the most blatantly anticompetitive forms of discrimination or blocking, and then limit the definition of specialized services to just those cases. For this situation to arise, a company must offer two communications services. One is a BIAS and subject to Open Internet rules. The other is not Internet access, and imposing Open Internet rules on this service would be inappropriate. There are cases like this, but they are limited.

For a service to be exempt from Open Internet discrimination and blocking rules without undue risk of the kind of oligopoly rents discussed in Section 1, that service should not be Internet access. Also, it should not be just another way to communicate with things that users would otherwise access over the Internet, as that would make it a trivial substitute for Internet. Finally, it should not operate over limited resources that would otherwise be used for Internet, as that would make it a simple Internet service with preferred access to shared resources.

There are services with the properties above for which the application of Open Internet rules would be inappropriate. This can occur when Open Internet rules apply to separate networks as if they were one network. Consider a company that offers two IP-based services. One gives a user access to the entire Internet, and the other gives a user access to a closed and highly secure network with endpoints that trust each other. For example, a company might use the closed network to connect its various offices and the residences of some trusted employees. The network for Internet access and the closed network are separate, in that traffic from one cannot reach the other, cannot carry malware to the other, and cannot even cause congestion to the other, which is precisely what makes the closed network valuable to its users. To further enhance the security and productivity of the closed network, traffic from some applications is blocked, which might violate BIAS Open Internet rules. Open Internet rules should not be imposed on the closed network, as this might allow traffic from unknown sources and unknown applications into the network, thereby reducing the network's value to users. Note that this is a true private network, and not a mere *virtual* private network (VPN) where private traffic and regular Internet traffic are intermixed within shared capacity. The VPN label should not automatically exempt a service from Open Internet rules.

Another example occurs when a company offers both an Internet and a telephone service that are entirely separate. Consider the case where the telephone network is upgraded from circuits to VOIP, while remaining separate from the network that provides Internet access. This alone should not make the telephone service subject to Open Internet rules, which would prohibit blocking non-VOIP packets. Otherwise, companies might be discouraged from improving technology. Both of the services above do not meet the definition of Internet access service and are not a trivial substitute for Internet because customers use them to access endpoints that are not on the open Internet. In the first case, those endpoints are other trusted computers within the company's private network. In the second case, the endpoint accessed is a VOIP-Public Switched Telephone Network (PSTN) gateway that is operated by the carrier and is accessible only through the carrier's closed network. In addition, central to the examples above, the service that is exempt from Open Internet rules is truly separate from Internet access.

Implicitly, some of the debate about specialized services is over what it means for two services to be separate. I propose here a simple litmus test that should be used to determine if services are separate, along with a condition to exclude services that are trivial substitutes for Internet.

A communications service can be considered a specialized service under Open Internet rules if (i) the primary use of the service is not to access content, services, or systems that are accessible through an Internet access service, and (ii) the service does not share capacity with Internet access.

Two services are said to “share capacity” if it is ever possible for utilization of one service to affect the performance of the other service.

The precise definition of sharing capacity is important. For example, when a telephone network offers a traditional DSL service, circuit-switched telephone and Internet services do operate over the same copper wire, but they do not share capacity. Internet and telephone traffic travel within separate frequency bands in the last-mile connection, arrive at separate switches at the central office, and are forwarded into separate nationwide networks (i.e., the PSTN and the Internet). If there is a high volume of Internet traffic on the last-mile link, in the IP router closest to the user, or in any link or router anywhere on the global Internet, this will not affect telephone service in any way. Similarly, a high volume of telephone traffic will not reduce the performance of the Internet service. This would not change if the telephone service was converted to VOIP, but the capacity allocated to telephone service is fixed, as is the capacity allocated to Internet service. In contrast, if VOIP packets are sent over the same last-mile link as Internet traffic and VOIP packets are simply given transmission priority over Internet traffic based on VLAN label, then telephone utilization would somewhat degrade the quality of Internet service. Contrary to what some ISPs will claim, this latter case is an example of shared capacity, and the proposed definition above makes this clear.

The impact of a specialized services rule on telephone and cable TV services deserves particular attention, in part because these two services have played a particularly important role in the business case for Internet access (e.g., as part of “Triple Play” service), and in part because these two services have their own regulations. Some IP-based telephone and cable TV services would qualify as specialized services under the proposed rule above, but some would not. For example, with AT&T’s U-verse service, a customer who begins watching a pay-per-view video may notice a sudden decrease in Internet performance because capacity is shared. AT&T adopted this architecture years ago, and if the FCC were to decide now that this technical approach makes the IP-based service used for video distribution subject to Open Internet rules, it would be disruptive. Consequently, I propose the following rule to specifically address IP-based telephone and cable TV services.

A communications service can be considered a specialized service under Open Internet rules if the service is only used to provide a service that is subject to telephone regulations or to cable TV regulations.

6.3 Summary

It is important that the FCC adopt appropriate rules for specialized services. Failure to do so could provide a loophole that would allow even the most harmful forms of discrimination to gain widespread use, or it could have the effect of denying Internet users access to valuable applications that require good quality of service. The following table summarizes the policy approaches in Sections 6.1 and 6.2.

	<i>Discrimination rules for BIAS</i>	<i>Definition of Specialized Services</i>	<i>Rules for Specialized Services</i>
Option 1	May or may not allow reasonable discrimination and the ability to charge for it.	Definition is broad to support all traffic streams that could benefit from superior QoS.	Unreasonable discrimination and blocking are prohibited.
Option 2	Allow reasonable discrimination and the ability to charge for better QoS. Prohibit unreasonable discrimination.	Definition is narrow, including only services that meet one of these conditions 1. The primary use of the service is not to access content, services, or systems that are accessible through an Internet access service, and (ii) the service does not share capacity with Internet access. <i>Or</i> 2. The service is only used to provide a service that is subject to telephone regulations or to cable TV regulations.	Few rules, if any (other than transparency).

References

- [CP14] B. A. Cherry & J. M. Peha, *The Telecom Act of 1996 Requires the FCC to Classify Commercial Internet Access as a Telecom Service*, Comments in the Matter of Protecting and Promoting the Open Internet, FEDERAL COMMUNICATIONS COMMISSION GN DOCKET NO. 14-28 (Dec. 22, 2014).
- [FC15] FEDERAL COMMUNICATIONS COMMISSION, *Protecting and Promoting the Open Internet*, Report and Order on Remand, Declaratory Ruling, and Order, GN Docket No. 14-28 (March 12, 2015).
- [FC17] FEDERAL COMMUNICATIONS COMMISSION, *Restoring Internet Freedom*, FEDERAL COMMUNICATIONS COMMISSION WC Docket No. 17-108 (May 23, 2017).
- [PE06] J. M. Peha, *The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy*, TELECOM POLICY RESEARCH CONFERENCE (TPRC) (2006), and INTERNATIONAL JOURNAL OF COMMUNICATION (2007).
- [PE14] J. M. Peha, A “Most Favored Nation Approach to an Open Internet,” Comments in the Matter of Protecting and Promoting the Open Internet, FEDERAL COMMUNICATIONS COMMISSION GN DOCKET NO. 14-28 (July 15, 2014).
- [PE15] J. M. Peha, Appropriate Rules for Managed or Specialized Services, Comments in the Matter of Protecting and Promoting the Open Internet, FEDERAL COMMUNICATIONS COMMISSION GN Docket No. 14-28 (Jan. 5, 2015).
- [PE16] J. M. Peha, The Network Neutrality Battles that will Follow Reclassification, *I/S: A Journal of Law and Policy for the Information Society* (2016).
- [PE17] J. M. Peha, Fallacies Behind Reclassifying Broadband Internet Access Service as an Information Service, *Comments in the Matter of Restoring Internet Freedom*, FEDERAL COMMUNICATIONS COMMISSION WC Docket No. 17-108 (July 17, 2017).
- [US14] United States Court of Appeals for District of Columbia Circuit, *Verizon v. FCC*, No. 11-1355, Jan. 14, 2014.